

Amendments to the Claims:

This listing of claims will replace all prior versions and listings of claims in the application. Added text is indicated by underlining, deleted text is indicated by ~~striking through~~. Changes are identified by a change bar at the left edge of text.

Listing of Claims:

1 1. (currently amended) A system for enabling queries to a database to be
2 processed comprising:
3 a database system;
4 an application system for providing SQL database queries to the database system,
5 the database system coupled to the application system via a first connection over a network;
6 a storage system having a shared volume to store results from SQL database
7 queries made to the database system;
8 a first data path to provide a data connection between the storage system and the
9 application system, wherein the application system can directly access query results on the
10 storage system without communicating via the first connection;
11 a second data path to provide a data connection between the storage system and
12 the database system, wherein the database system directly stores query results to the storage
13 system via the second data path without communicating via the first connection; and
14 a return path selector coupled to the database system for selecting a return path
15 over which to return the results from queries made to the database system, the return path
16 selector selecting from among at least the first connection over the network or the first data path
17 between the storage system and the application system, wherein the return path selector
18 determines a data path based upon one or more attributes of the query results;
19 wherein, when the return path is chosen to be the storage system, the results are
20 sent to the storage system as a file and an address in the storage system for the file is provided to
21 the application system using the first connection; and

22 wherein the file has associated therewith a key and the key is used to control
23 access to the results, and also has associated therewith a flag to indicate status of the file, and
24 wherein the flag indicates at least one of whether the file is being written, is ready to be read, is
25 being read, and is available to be deleted.

1 2. (currently amended) A system as in claim 1 further comprising a request
2 path selector coupled to the application system for selecting a request path over which to send
3 query data comprising the SQL database queries made to the database system, the request path
4 selector selecting from among at least the first connection or the storage system, wherein the key
5 and flag are stored in the storage system.

1 3. (original) A system as in claim 1 wherein the storage system is coupled to
2 each of the application system and the database system using a switch.

1 4. (original) A system as in claim 3 wherein a database hub system is used
2 to couple the application system and the database system.

1 5. (currently amended) A system as in claim 1 wherein the results from the
2 query have a size, and the return path selector chooses a return path based on the size of the
3 results, and wherein said return selector compares the size of the results to a threshold to choose
4 said result path, and said threshold is set based on a current workload of the LAN.

1 6. (currently amended) A system as in claim 5 wherein the return path
2 selector chooses a return path based on a prediction of the size of the results, and wherein said
3 prediction is based on similar queries from a query execution history.

1 7. (original) A system as in claim 1 wherein the return path selector chooses
2 a return path based on a measurement of throughput of the first connection.

1 8. (canceled)

1 9. (original) A system as in claim 1 wherein after the results are used by the
2 application system, the application system designates the results as used, thereby enabling them
3 to be erased from the storage system at a later time.

1 10. (original) A system as in claim 1 wherein after the query data is used by
2 the database system, the database system reuses the query data for a further query.

1 11. (canceled)

1 12. (canceled)

1 13. (canceled)

1 14. (previously presented) A system as in claim 10 wherein after the query
2 data is used by the database system, the database system designates the query data as used,
3 thereby enabling them to be erased from the storage system at a later time.

1 15. (canceled)

1 16. (previously presented) A system as in claim 1 further comprising a hub
2 system coupled to each of the application system and the database system.

1 17. (currently amended) A system for enabling queries to a database to be
2 processed comprising:
3 a database system;
4 an application system for providing SQL database queries to the database system
5 coupled to the application system via a first connection over a network;
6 a storage system having a volume to store results from SQL database queries
7 made to the database system;

a first data path to provide a data connection between the storage system and the application system, wherein the application system can directly access query results on the storage system without communicating via the first connection;

a second data path to provide a data connection between the storage system and the database system, wherein the database system directly stores query results to the storage system via the second data path without communicating via the first connection; and

a request path selector coupled to the application system for selecting a request path over which to send query data comprising the SQL database queries made to the database system, the request path selector selecting from among at least the first connection over the network or the first data path between the storage system and the application system, wherein the request path selector determines a data path based upon one or more attributes of the query data;

wherein, when the request path is chosen to be the storage system, the query data are sent to the storage system as a file and an address in the storage system for the file is provided to the database system using the first connection; and

wherein the file has associated therewith a key and the key is used to control access to the results, and also has associated therewith a flag to indicate status of the file, and the flag indicates at least one of whether the file is being written, is ready to be read, is being read, and is available to be deleted.

18. (original) A system as in claim 17 wherein the storage system is coupled to each of the application system and the database system using a switch.

19. (original) A system as in claim 18 wherein a database hub system is used to couple the application system and the database system.

20. (currently amended) A system as in claim 17 wherein the query data have a size, and the request path selector chooses a request path based on the size of the query data, and wherein if there is not enough space in said storage system, garbage collection is performed before storing said query data to said storage system.

1 21. (canceled)

1 22. (canceled)

1 23. (canceled)

1 24. (canceled)

1 25. (original) A system as in claim 17 further comprising a return path
2 selector coupled to the database system for selecting a return path over which to return results
3 from queries made to the database system, the return path selector selecting from among at least
4 the first connection or the storage system.

1 26. (currently amended) A system as in claim 25 wherein the results from the
2 query have a size, and the return path selector chooses a return path based on the size of the
3 results, and wherein if there is not enough space in said storage system, garbage collection is
4 performed before storing said query data to said storage system.

1 27. (previously presented) A system as in claim 17 wherein when the return
2 path is chosen to be the storage system, the results are sent to the storage system as a file and an
3 address in the storage system for the file is provided to the application system using the first
4 connection.

1 28. (currently amended) A system for enabling queries to a database to be
2 processed comprising:
3 a database system;
4 an application system for providing SQL database queries to the database system
5 coupled to the application system via a first connection over a network, the application system
6 including a database access system, and the database system including a gateway system;
7 a storage system having a shared volume to store results from SQL database
8 queries made to the database system;

a first data path to provide a data connection between the storage system and the application system, wherein the application system can directly access query results on the storage system without communicating via the first connection;

a second data path to provide a data connection between the storage system and the database system, wherein the database system directly stores query results to the storage system via the second data path without communicating via the first connection; and

the gateway system including a return path selector for selecting a return path over which to return the results from queries made to the database system, the return path selector selecting from among at least the first connection over the network or the first data path between the storage system and the application system, wherein the return path selector determines a data path based upon one or more attributes of the query results;

wherein, when the request path is chosen to be the storage system, the query data are sent to the storage system as a file and an address in the storage system for the file is provided to the database system using the first connection; and

wherein the file has associated therewith a key and the key is used to control access to the results, and also has associated therewith a flag to indicate status of the file, and the flag indicates at least one of whether the file is being written, is ready to be read, is being read, and is available to be deleted.

29. (previously presented) A system as in claim 28 wherein the gateway system includes a request path selector for selecting a request path over which to send query data comprising the SQL database queries made to the database system, the request path selector selecting from among at least the first connection or the storage system.

30. (currently amended) A system for enabling queries to a database to be processed comprising:

a database system;

an application system for providing SQL database queries to the database system coupled to the application system via a first connection over a network, the application system including a database access system, and the database system including a gateway system;

a storage system having a volume to store results from SQL database queries made to the database system;

a first data path to provide a data connection between the storage system and the application system, wherein the application system can directly access query results on the storage system without communicating via the first connection;

a second data path to provide a data connection between the storage system and the database system, wherein the database system directly stores query results to the storage system via the second data path without communicating via the first connection; and

the database access system including a request path selector for selecting a request path over which to send data comprising the SQL database queries made to the database system, the request path selector selecting from among at least the first connection over the network or the first data path between the application system and the storage system;

wherein, when the request path is chosen to be the storage system, the query data are sent to the storage system as a file and an address in the storage system for the file is provided to the database system using the first connection; and

wherein the file has associated therewith a key and the key is used to control access to the results, and also has associated therewith a flag to indicate status of the file, and the flag indicates at least one of whether the file is being written, is ready to be read, is being read, and is available to be deleted.

31. (currently amended) A system for enabling queries to a database to be processed comprising:

a database system;

an application system for providing SQL database queries to the database system, the database system coupled to the application system via a communications network connection;

7 a switch coupled to each of the database system and the application system;
8 a storage system coupled to the switch, the storage system having a volume to
9 store results from SQL database queries made to the database system;
10 a first data path to provide a data connection between the storage system and the
11 application system, wherein the application system can directly access query results on the
12 storage system without communicating via the first connection;

13 a second data path to provide a data connection between the storage system and
14 the database system, wherein the database system directly stores query results to the storage
15 system via the second data path without communicating via the first connection; and

16 a return path selector coupled to the database system for selecting a return path
17 over which to return the results from queries made to the database system, the return path
18 selector selecting from among at least the first connection over the network or the first data path
19 between the storage system and the application system, wherein the return path selector
20 determines a data path based upon one or more attributes of the query results;

21 wherein the database system is configured to generate a key for the results to
22 identify their location, send the key over the network to the query provider, and encrypt at least
23 one of the key and the results.

1 32. (currently amended) A system as in claim 31 further comprising a request
2 path selector coupled to the application system for selecting a request path over which to send
3 query data comprising the SQL database queries made to the database system, the request path
4 selector selecting from among at least the communications network connection or the switch, and
5 wherein said return selector compares the size of the results to a threshold to choose said result
6 path, and said threshold is set based on a current workload of the LAN.

1 33. (currently amended) In a system having a query provider which provides
2 SQL database queries to a database system connected to the query provider by a first connection
3 over a network, the query provider and the database system being each coupled to a storage

4 system via different paths, a method of returning results to the query provider, a method
5 comprising: ~~<carriage-return>~~
6 storing results from SQL database queries made to the database system in the
7 storage system as a file at an address which can be accessed separately by the query provider via
8 a second connection independent of the network and by the database system via a third
9 connection independent of the network; and ~~<carriage-return>~~
10 sending the address of the results in the storage system via the first connection
11 over the network to the query provider.

1 34. (original) A method as in claim 33 further comprising:
2 generating a key for the results to identify their location; and
3 sending the key over the network to the query provider.

1 35. (original) A method as in claim 34 further comprising a step of, at the
2 query provider, retrieving the results from the storage system.

1 36. (original) A method as in claim 34 further comprising encrypting at least
2 one of the key and the results.

1 37. (previously presented) A method as in claim 33 wherein the query
2 provider provides query data comprising the SQL database queries to the database by storing the
3 query data in the storage system at a location and sending information about the location over the
4 first connection to the database system.

1 38. (currently amended) A method as in claim ~~28~~33 further comprising, at
2 the database system, the steps of:
3 retrieving the query data from the storage system; and
4 using the query data to obtain the results.

1 39. (original) A method as in claim 33 further comprising providing a flag
2 associated with the results to indicate whether the results are ready to be read by the query
3 provider.

1 40. (original) A method as in claim 37 further comprising providing a flag
2 associated with the results to indicate whether the results have been read by the query provider.

1 41. (currently amended) In a data storage system connected to an application
2 system and a database system via a network, a method comprising:

3 receiving from the database system over the network, results of execution of SQL
4 database queries, the SQL database queries being sent to the database system by the application
5 system;

6 storing the results of execution of SQL database queries in a storage area that the
7 database system and the application system can access separately via separate connections to the
8 storage area, the separate connections comprising:

9 a first data path to provide a data connection between the storage system
10 and the application system, wherein the application system can directly access query results on
11 the storage system without communicating via the first connection; and

12 a second data path to provide a data connection between the storage
13 system and the database system, wherein the database system directly stores query results to the
14 storage system via the second data path without communicating via the first connection;

15 determining a return path for the results of the execution of queries to the
16 application system based upon one or more attributes of the results of the execution of queries;
17 and

18 sending, in response to a request from the application system, the results of
19 execution of queries to the application system over the network if the network is determined to
20 be the return path or returning an address in the shared volume for the results of the execution of
21 the query if the shared volume is determined to be the return path;

further comprising:
generating a key for the results to identify their location;
sending the key over the network to the query provider; and
encrypting at least one of the key and the results.

42. (currently amended) In a system having an application system, a database system connected to the application system via a first network connection, a return path selector coupled to the database system for selecting a return path over which to return the results from queries made to the database system, and a data storage system connected to the application system via a first data path and connected to the database system through a second data path, the first data path providing a data connection between the data storage system and the application system through which the application system can directly access query results on the data storage system without communicating via the first connection, and the second data path providing a data connection between the data storage system and the database system through which the database system directly stores query results to the data storage system via the second data path without communicating via the first connection, a method, comprising the steps of:

sending an SQL database query from the application system to the database system by using the first network connection;

selecting a return path over which to return the results from SQL database queries made to the database system from among at least the first network connection or the first data path between the storage system and the application system, wherein the return path selector determines the return path based upon one or more attributes of the query results;

storing a result of execution of the query in a shared volume of the data storage system that can be accessed by the application system via the first data path and by the database system via the second data path; and

obtaining at the application system, the result of execution of the query from the storage system via the first data path without going through the first connection over the network;

further comprising:

25 | generating a key for the results to identify their location;
26 | sending the key over the network to the query provider; and
27 | encrypting at least one of the key and the results.

1 43. (original) The method of claim 42, wherein the first connection is a Local
2 Area Network, and wherein the second connection is a Storage Area Network.